

Closed-form EM for Sparse Coding And its Application to Source Separation

Jörg Lücke* and Abdul-Saboor Sheikh*

FIAS, Goethe-University Frankfurt, 60438 Frankfurt, Germany
{luecke, sheikh}@fias.uni-frankfurt.de

Abstract. We define and discuss the first sparse coding algorithm based on closed-form EM updates and continuous latent variables. The underlying generative model consists of a standard ‘spike-and-slab’ prior and a Gaussian noise model. Closed-form solutions for E- and M-step equations are derived by generalizing probabilistic PCA. The resulting EM algorithm can take all modes of a potentially multi-modal posterior into account. The computational cost of the algorithm scales exponentially with the number of hidden dimensions. However, with current computational resources, it is still possible to efficiently learn model parameters for medium-scale problems. Thus the model can be applied to the typical range of source separation tasks. In numerical experiments on artificial data we verify likelihood maximization and show that the derived algorithm recovers the sparse directions of standard sparse coding distributions. On source separation benchmarks comprised of realistic data we show that the algorithm is competitive with other recent methods.

1 Introduction

Probabilistic generative models are a standard approach to model data distributions and to infer instructive information about the data generating process. Methods like principle component analysis, factor analysis, or sparse coding (SC) (e.g., [14]) have all been formulated in the form of probabilistic generative models. Moreover, independent component analysis (ICA), which is a very popular approach to blind source separation, can also be recovered from sparse coding in the limit of zero observation noise (e.g., [4]).

A standard procedure to optimize parameters in generative models is the application of Expectation Maximization (EM) (e.g., [13]). However, for many generative models the optimization using EM is analytically intractable. For stationary data only the most elementary models such as mixture models and factor analysis (which contains probabilistic PCA as special case) have closed-form solutions for E- and M-step equations. EM for more elaborate models requires approximations. In particular, sparse coding models ([14,9,16] and many more) require approximations because integrals over the latent variables do not have closed-form solutions.

* joint first authorship

In this work we study a generative model that combines the Gaussian prior of probabilistic PCA (p-PCA) with a binary prior distribution. Distributions combining binary and continuous parts have been discussed and used as priors before (e.g., [11]) and are commonly referred to as ‘spike-and-slab’ distributions. Also sparse coding variants with spike-and-slab distributions have been studied previously (compare [20,7,18,15,8,12]). However, in this work we show that combining binary and Gaussian latents maintains the p-PCA property of having a closed-form solution for EM optimization. We can, therefore, derive an algorithm that uses exact posteriors with potentially many modes to update model parameters.




Fig. 1. Distributions generated by the GSC generative model. The left column shows the distributions generated for $\pi_h = 1$ for all h . In this case the model generates p-PCA distributions. The middle column shows an intermediate value of π_h . The generated distributions are not Gaussians anymore but have a slight star shape. The right column shows distributions for small values of π_h . The generated distributions have a salient star shape similar to standard sparse coding distributions.

2 The Gaussian Sparse Coding (GSC) model

Let us first consider a pair of H -dimensional i.i.d. latent vectors, a continuous $\mathbf{z} \in \mathbb{R}^H$ and a binary $\mathbf{s} \in \{0, 1\}^H$ with:

$$p(\mathbf{s} | \Theta) = \prod_{h=1}^H \pi_h^{s_h} (1 - \pi_h)^{1-s_h} = \text{Bernoulli}(\mathbf{s}; \boldsymbol{\pi}) \text{ and } p(\mathbf{z} | \Theta) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbb{I}_H), \quad (1)$$

where π_h parameterizes the probability of non-zero entries. After generation, both hidden vectors are combined using a pointwise multiplication operator: i.e., $(\mathbf{s} \odot \mathbf{z})_h = s_h z_h$ for all h . The resulting hidden random variable is a vector of continuous values and zeroes, and it follows a ‘spike-and-slab’ distribution. Given a hidden vector (which we will denote by $\mathbf{s} \odot \mathbf{z}$), we generate a D -dimensional observation $\mathbf{y} \in \mathbb{R}^D$ by linearly combining a set of basis functions \mathbf{W} and adding Gaussian noise:

$$p(\mathbf{y} | \mathbf{s}, \mathbf{z}, \Theta) = \mathcal{N}(\mathbf{y}; \mathbf{W}(\mathbf{s} \odot \mathbf{z}), \Sigma), \quad (2)$$

where $\mathbf{W} \in \mathbb{R}^{D \times H}$ is the matrix containing the basis functions \mathbf{W}_h as columns, and $\Sigma \in \mathbb{R}^{D \times D}$ is a covariance matrix parameterizing the data noise. The latents’ priors (1) together with their pointwise combination and the noise distribution (2) define the generative model under consideration. As a special case, the model contains probabilistic PCA (or factor analysis). This can easily be seen by setting all π_h equal to one.

The model (1) to (2) is capable of generating a broad range of distributions including sparse coding like distributions. This is illustrated in Fig. 1 where the parameters π_h allow for continuously changing PCA-like to a SC-like distribution.

While the generative model itself has been studied previously [20,18,15,8], we will show that a closed-form EM algorithm can be derived, which can be applied to blind source separation tasks. We will refer to the generative model (1) to (2) as the *Gaussian Sparse Coding* (GSC) model in order to stress that a specific spike-and-slab prior (Gaussian slab) in conjunction with a Gaussian noise model is used. The GSC model is thus an instance of the spike-and-slab sparse coding model (or alternatively known *sparse factor analysis* models; see e.g., [20,18,15,8]).

2.1 Expectation Maximization (EM) for Parameter Optimization

Consider a set of N independent data points $\{\mathbf{y}^{(n)}\}_{n=1,\dots,N}$ with $\mathbf{y}^{(n)} \in \mathbb{R}^D$. For these data we seek parameters $\Theta = (W, \Sigma, \boldsymbol{\pi})$ that maximize the data likelihood $\mathcal{L} = \prod_{n=1}^N p(\mathbf{y}^{(n)} | \Theta)$ under the GSC generative model. We employ Expectation Maximization (EM) algorithm for parameter optimization. The EM algorithm [13] optimizes the data likelihood w.r.t. the parameters Θ by iteratively maximizing the free-energy given by:

$$\mathcal{F}(\Theta^{\text{old}}, \Theta) = \sum_{n=1}^N \sum_{\mathbf{s}} \int_{\mathbf{z}} p(\mathbf{s}, \mathbf{z} | \mathbf{y}^{(n)}, \Theta^{\text{old}}) \left[\log(p(\mathbf{y}^{(n)} | \mathbf{s}, \mathbf{z}, \Theta)) + \log(p(\mathbf{s} | \Theta)) + \log(p(\mathbf{z} | \Theta)) \right] d\mathbf{z} + H(\Theta^{\text{old}}), \quad (3)$$

where $H(\Theta^{\text{old}})$ is an entropy term only depending on parameter values held fixed during the optimization of \mathcal{F} w.r.t. Θ . Note that integration over the hidden space involves an integral over the continuous part and a sum over the binary part.

Optimizing the free-energy consists of two steps: given the current parameters Θ^{old} the posterior probability is computed in the E-step; and given the posterior, $\mathcal{F}(\Theta^{\text{old}}, \Theta)$ is maximized w.r.t. Θ in the M-step. Iteratively applying E- and M-steps locally maximizes the data likelihood.

M-step parameter updates: Let us first consider the maximization of the free-energy in the M-step before considering expectation values w.r.t. to the posterior in the E-step. Given a generative model, conditions for a maximum free-energy are canonically derived by setting the derivatives of $\mathcal{F}(\Theta^{\text{old}}, \Theta)$ w.r.t. the second argument to zero. For the GSC model we obtain the following parameter updates:

$$W = \left(\sum_{n=1}^N \mathbf{y}^{(n)} \langle \mathbf{s} \odot \mathbf{z} \rangle_n^T \right) \left(\sum_{n=1}^N \langle (\mathbf{s} \odot \mathbf{z})(\mathbf{s} \odot \mathbf{z})^T \rangle_n \right)^{-1}, \quad (4)$$

$$\Sigma = \frac{1}{N} \sum_{n=1}^N \left[\mathbf{y}^{(n)} (\mathbf{y}^{(n)})^T - 2(W \langle \mathbf{s} \odot \mathbf{z} \rangle_n) (\mathbf{y}^{(n)})^T + W \langle (\mathbf{s} \odot \mathbf{z})(\mathbf{s} \odot \mathbf{z})^T \rangle_n W^T \right]$$

$$\text{and } \boldsymbol{\pi} = \frac{1}{N} \sum_{n=1}^N \langle \mathbf{s} \rangle_n, \text{ where } \langle f(\mathbf{s}, \mathbf{z}) \rangle_n = \sum_{\mathbf{s}} \int_{\mathbf{z}} p(\mathbf{s}, \mathbf{z} | \mathbf{y}^{(n)}, \Theta^{\text{old}}) f(\mathbf{s}, \mathbf{z}) d\mathbf{z}. \quad (5)$$

Equations (4) to (5) define a new set of parameter values $\Theta = (W, \Sigma, \pi)$ given the current values Θ^{old} . These 'old' parameters are only used to compute the sufficient statistics $\langle \mathbf{s} \rangle_n$, $\langle \mathbf{s} \odot \mathbf{z} \rangle_n$ and $\langle (\mathbf{s} \odot \mathbf{z})(\mathbf{s} \odot \mathbf{z})^T \rangle_n$ of the model.

Expectation Values: Although the derivation of M-step equations can be analytically intricate, it is the E-step that, for most generative models, poses the major challenge. Source of the problems involved are analytically intractable integrals required for posterior distributions and for expectation values w.r.t. the posterior. The true posterior is therefore often replaced by an approximate distribution (see, e.g., [2,16]) or in the form of factored variational distributions [6,5]. The most frequently used approximation is the maximum-a-posterior (MAP) estimate (see, e.g., [14,9]) which replaces the true posterior by a delta-function around the posterior's maximum value. Alternatively, analytically intractable expectation values are often approximated using sampling approaches. Using approximations always implies, however, that many analytical properties of exact EM are not maintained. Approximate EM iterations may, for instance, decrease the likelihood or may not recover (local or global) likelihood optima in many cases. There are nevertheless, a limited number of models with exact EM solutions; e.g., mixture models such as the mixture-of-Gaussians, p-PCA or factor analysis etc. Our novel work here extends the set of known models with exact EM solutions. By following along the same lines as for the p-PCA derivations, we maintain in our E-step the analytical tractability of computing expectation values w.r.t. the posterior of the GSC model (5).

Posterior Probability: First observe that the discrete latent variable \mathbf{s} of the GSC model can be directly combined with the basis functions, i.e., $W(\mathbf{s} \odot \mathbf{z}) = \tilde{W}_s \mathbf{z}$, where $(\tilde{W}_s)_{dh} = W_{dh} s_h$. Now we apply the Bayes' rule to write down the posterior:

$$p(\mathbf{s}, \mathbf{z} | \mathbf{y}^{(n)}, \Theta) = \frac{\mathcal{N}(\mathbf{y}^{(n)}; \tilde{W}_s \mathbf{z}, \Sigma) \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbb{1}_H) p(\mathbf{s} | \Theta)}{\sum_{\mathbf{s}'} \int \mathcal{N}(\mathbf{y}^{(n)}; \tilde{W}_{s'} \mathbf{z}', \Sigma) \mathcal{N}(\mathbf{z}'; \mathbf{0}, \mathbb{1}_H) p(\mathbf{s}' | \Theta) d\mathbf{z}'}. \quad (6)$$

Note that given a state \mathbf{s} in (6), the Gaussian governing the observations $\mathbf{y}^{(n)}$ is only dependent on the Gaussian over the continuous latent \mathbf{z} , which is analytically independent of \mathbf{s} . We can exploit this joint relation to refactorize the Gaussians. Using Gaussian identities the posterior can be rewritten as:

$$\begin{aligned} p(\mathbf{s}, \mathbf{z} | \mathbf{y}^{(n)}, \Theta) &= \frac{\mathcal{N}(\mathbf{y}^{(n)}; \mathbf{0}, C_s) p(\mathbf{s} | \Theta) \mathcal{N}(\mathbf{z}; \boldsymbol{\kappa}_s^{(n)}, \Lambda_s)}{\sum_{\mathbf{s}'} \mathcal{N}(\mathbf{y}^{(n)}; \mathbf{0}, C_{s'}) p(\mathbf{s}' | \Theta) \int \mathcal{N}(\mathbf{z}'; \boldsymbol{\kappa}_{s'}^{(n)}, \Lambda_{s'}) d\mathbf{z}'} \\ &= p(\mathbf{s} | \mathbf{y}^{(n)}, \Theta) \mathcal{N}(\mathbf{z}; \boldsymbol{\kappa}_s^{(n)}, \Lambda_s), \end{aligned} \quad (7)$$

$$\begin{aligned} \text{where } C_s &= \tilde{W}_s \tilde{W}_s^T + \Sigma, \quad \Lambda_s = (\tilde{W}_s^T \Sigma^{-1} \tilde{W}_s + \mathbb{1}_H)^{-1} \\ \text{and } \boldsymbol{\kappa}_s^{(n)} &= \Lambda_s \tilde{W}_s^T \Sigma^{-1} \mathbf{y}^{(n)}. \end{aligned} \quad (8)$$

Equations (7) to (8) represent the crucial result for the computation of the E-step below because, first, they show that the posterior does not involve analytically intractable integrals and, second, for fixed \mathbf{s} and $\mathbf{y}^{(n)}$ the dependency on \mathbf{z} follows a Gaussian distribution. This special form allows for the derivation of analytical expressions for

the expectation values as required for the M-step parameter updates.

E-step Equations: Derived from (7), the expectation values are computed as:

$$\langle s \rangle_n = \sum_s p(s | \mathbf{y}^{(n)}, \Theta) s, \quad \langle s \odot \mathbf{z} \rangle_n = \sum_s p(s | \mathbf{y}^{(n)}, \Theta) \kappa_s^{(n)} \quad (9)$$

$$\text{and } \langle (s \odot \mathbf{z})(s \odot \mathbf{z})^T \rangle_n = \sum_s p(s | \mathbf{y}^{(n)}, \Theta) (\Lambda_s + \kappa_s^{(n)}(\kappa_s^{(n)})^T). \quad (10)$$

Note that we have to use the current values $\Theta = \Theta^{\text{old}}$ for all parameters on the right-hand-side. The E-step equations (9) to (10) represent a closed-form solution for expectation values required for the closed-form M-step (4) to (5).

Relation to the Mixture of Gaussians: The special form of the posterior in (7) allows the derivation of a closed-form expression of the data likelihood: i.e., $p(\mathbf{y} | \Theta) = \sum_s \text{Bernoulli}(s; \pi) \mathcal{N}(\mathbf{y}^{(n)}; \mathbf{0}, C_s)$. Note that in principle, this form can be reproduced by a Gaussian mixture model. However, such a model would consist of 2^H mixture components, with strongly dependent mixing proportions and covariance matrices C_s . Closed-form EM-updates can in general not be derived for such dependencies. The standard updates for mixtures of Gaussians require independently parameterized mixing proportions and components. Therefore, the closed-form EM-solutions for the GSC model is not a consequence of closed-form EM for classical Gaussian mixtures.

3 Numerical Experiments

GSC parameter optimization is non-convex. However, as for all algorithms based on closed-form EM, the GSC algorithm always increases the data likelihood at least to a local maxima. We first numerically investigate how frequently local optima are obtained. Later we assess model's performance on more practical tasks.

Model verification: First, we verify on artificial data that the algorithm increases the likelihood and that it can recover the parameters of the generating distribution. For this, we generated $N = 500$ data points $\mathbf{y}^{(n)}$ from the GSC generative model (1) to (2) with $D = H = 2$. We used randomly initialized generative parameters¹. The algorithm was run 250 times on the generated data. For each run we performed 300 EM iterations. For each run, we randomly and uniformly initialized π_h between 0.05 and 10, set Σ to the covariance across the data points, and the elements of W we chose to be independently drawn from a normal distribution with zero mean and unit variance. In all runs the generating parameter values were recovered with high accuracy. Runs with different generating parameters produced essentially the same results.

Recovery of sparse directions: To test the model's robustness w.r.t. a relaxation of the GSC assumptions, we applied the GSC algorithm to data generated by standard sparse coding models. We used a standard Cauchy prior and a Gaussian noise model [14] for

¹ We obtained W^{gen} by independently drawing each matrix entry from a normal distribution with zero mean and standard deviation 3. π_h^{gen} values were drawn from a uniform distribution between 0.05 and 1, $\Sigma = \sigma^{\text{gen}} \mathbb{1}_D$ (where σ^{gen} was uniformly drawn between 0.05 and 10).




Fig. 2. Histogram of likelihood values for 100 runs of the GSC algorithm on data generated by a SC model with Cauchy prior. Almost all runs converged to high likelihood values.

data generation. Fig. 3 second panel shows data generated by this sparse coding model while the first panel shows the prior density along one of its hidden dimensions. We generated $N = 500$ data points with $H = D = 2$. We then applied the GSC algorithm with the same parameter initialization as in the previous experiment. We performed 100 trials using 300 EM iterations per trial. Again, the algorithm converged to high likelihood values in most runs (see Fig. 2). As a performance measure for this experiment we investigated how well the heavy tails (i.e., the sparse directions) of standard SC were recovered. As a performance metric, we used the Amari index [1]:




Fig. 3. Comparison of standard sparse coding and GSC. **Left panels:** Cauchy distribution (along one hidden dimension) as a standard SC prior [14] and data generated by it. **Right panels:** Spike-and-slab distribution (one of the hidden dimensions) inferred by the GSC algorithm along with inferred sparse directions (solid red lines) and posterior data density contours (dotted red lines).

$$A(W) = \frac{1}{2H(H-1)} \sum_{h,h'=1}^H \left(\frac{|O_{hh'}|}{\max_{h''} |O_{hh''}|} + \frac{|O_{hh'}|}{\max_{h''} |O_{h''h'}|} \right) - \frac{1}{H-1} \quad (11)$$

where $O_{hh'} := (W^{-1}W^{\text{gen}})_{hh'}$. The mean Amari index of all runs with high likelihood values was below 10^{-2} , which shows a very accurate recovery of the sparse directions. Fig. 3 (right panel) visualizes the distribution recovered by the GSC algorithm in a typical run. The dotted red lines show the density contours of the learned distribution $p(\mathbf{y} | \Theta)$. High accuracy in the recovery of the generating sparse directions (solid black lines) can be observed by comparison with the recovered directions (solid red lines). The results of experiments are qualitatively the same if we increase the number of hidden and observed dimensions; e.g., for $H = D = 4$ we found the algorithm converged to a high likelihood in 91 (with average Amari index below 10^{-2}) of 100 runs.

Other than standard SC with Cauchy prior, we also ran the algorithm on data generated by SC with Laplace prior [14,9]. There for $H = D = 2$, we converged to high likelihood values in 99 of 100 runs with an average Amari index 0.06. In the experiment

Table 1. Performance of different algorithms on benchmarks for source separation. Data for NG-LICA, KICA, FICA, and JADE are taken from [17]. Performances are compared based on the Amari index (11). Bold values highlight the best performing algorithm(s).

datasets		Amari index (standard deviation)					
name	N	GSC	GSC^\perp	NG-LICA	KICA	FICA	JADE
10halo	200	0.34(0.05)	0.29(0.03)	0.29(0.02)	0.38(0.03)	0.33(0.07)	0.36(0.00)
	500	0.27(0.01)	0.27(0.01)	0.22(0.02)	0.37(0.03)	0.22(0.03)	0.28(0.00)
Sergio7	200	0.23(0.06)	0.20(0.06)	0.04(0.01)	0.38(0.04)	0.05(0.02)	0.07(0.00)
	500	0.18(0.05)	0.17(0.03)	0.05(0.02)	0.37(0.03)	0.04(0.01)	0.04(0.00)
Speech4	200	0.25(0.05)	0.17(0.04)	0.18(0.03)	0.29(0.05)	0.20(0.03)	0.22(0.00)
	500	0.11(0.04)	0.05(0.01)	0.07(0.00)	0.10(0.04)	0.10(0.04)	0.06(0.00)
c5signals	200	0.39(0.03)	0.44(0.05)	0.12(0.01)	0.25(0.15)	0.10(0.02)	0.12(0.00)
	500	0.41(0.05)	0.44(0.04)	0.06(0.04)	0.07(0.06)	0.04(0.02)	0.07(0.00)

with $H = D = 4$ the algorithm converged to a high likelihood in 97 of 100 runs. The average Amari index of all runs with high likelihoods was 0.07 in this case.

Source separation: We applied the GSC algorithm to publicly available benchmarks. We used the non-artificial benchmarks of [17]. The datasets mainly contain acoustic data obtained from ICALAB [3]. We generated the observed data by mixing the benchmark sources using randomly generated orthogonal mixing matrix (we followed [17]). Again, the Amari index (11) was used as a performance measure.




Fig. 4. Histogram of the deviation from orthogonality of the W matrix for 100 runs of the GSC algorithm on the Speech4 benchmark ($N = 500$). A clear cluster of the most orthogonal runs can automatically be detected: the threshold of runs considered is defined to be the minimum after the cluster (black arrow).

For all the benchmarks we used $N = 200$ and $N = 500$ data points (as selected by [17]). We applied GSC to the data using the same initialization as described before. For each experiment we performed 100 trials with a random new parameter initialization per trial. The first column of Tab. 1 list average Amari indices obtained including all trials per experiment². It is important to note that all the other algorithms listed in the comparison assume orthogonal mixing matrices, while the GSC algorithm does not. Therefore in the column ' GSC^\perp ' in Tab. 1, we report statistics that are only computed over the runs which inferred the most orthogonal W matrices. As a measure of orthogonality we used the maximal deviation from 90° between any two axes. Fig. 4 shows as an example a histogram of the maximal deviations of all trials on the Speech4 data with $N = 500$. As can be observed, we obtained a clear cluster of runs with high orthogonality. We observed worst performance of the GSC algorithm on the c5signals

² We obtained the reported results by diagonalizing the updated Σ in the M-step by setting $\Sigma = \sigma^2 \mathbb{1}_D$, where $\sigma^2 = \text{Tr}(\Sigma)/D$.

dataset. However, the dataset contains sub-Gaussian sources which in general can not be recovered by sparse coding approaches.

4 Discussion

The GSC algorithm falls into the class of standard SC algorithms. However, instead of a heavy-tail prior, it uses a spike-and-slab distribution. The algorithm has a distinguishing capability of taking the whole (potentially a multimodal) posterior into account for parameter optimization, which is in contrast to the MAP approximation of the posterior, which is a widely used approach for training SC models (see, e.g., [9,10]). Various sophisticated methods have been proposed to efficiently find the MAP (e.g., [19]). MAP based optimizations usually also require regularization parameters that have to be inferred (e.g., by means of cross-validation). Other approximations that take more aspects of the posterior into account are also being investigated actively (e.g., [16,12]). However, approximations can introduce learning biases. For instance, MAP and Laplace approximations assume monomodality in posterior estimation, which is not always the case.

Closed-form EM learning of the GSC algorithm also comes with a cost, which is exponential w.r.t. the number of hidden dimensions H . This can be seen by considering (8) where the partition function requires a summation over all binary vectors s (similar for expectation values w.r.t. the posterior). Nevertheless, we show in numerical experiments that the algorithm is well applicable to the typical range of source separation tasks. In such domains the GSC algorithm can benefit from taking potentially a multimodal posterior into account and inferring a whole set of model parameters including the sparsity per latent dimension. For instance, when using a number of hidden dimensions larger than the number of actual sources, the model can discard dimensions by setting π_h parameters to zero. The studied model could thus be considered as treating parameter inference in a more Bayesian way than, e.g., SC with MAP estimates (compare [9]). The second line of research aims at a fully Bayesian description of sparse coding and emphasises a large flexibility using estimations of the number of hidden dimensions and by being applicable with a range of different noise models. The great challenge of these general models is the procedure of parameter estimation. For the model in [12], for instance, the Bayesian methodology involves conjugate priors and hyperparameters in combination with Laplace approximation and different sampling schemes. While the aim, e.g., in [20,7,18,15,8,12] is a large flexibility, we aim at a generalization of sparse coding that maintains the closed-form EM solutions.

To summarize, we have studied a novel sparse coding algorithm and have shown its competitiveness on source separation benchmarks. Along with the reported results on source separation, the main contribution of this work is the derivation and numerical investigation of the (to the knowledge of the authors) first closed-form, exact EM algorithm for spike-and-slab sparse coding.

Acknowledgement. J. Lücke is funded by the German Research Foundation (DFG), grant LU 1196/4-1; A.-S. Sheikh is funded by the German BMBF, grant 01GQ0840.

References

1. S. Amari, A. Cichocki, and H. H. Yang. A new learning algorithm for blind signal separation. In *NIPS*, pages 757–763. MIT Press, 1996.
2. C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
3. A. Cichocki, S. Amari, K. Siwek, T. Tanaka, A. Phan, and R. Zdunek. ICALAB-MATLAB Toolbox Version 3, 2007.
4. P. Dayan and L. F. Abbott. *Theoretical Neuroscience*. MIT Press, Cambridge, 2001.
5. T. Jaakkola. Tutorial on variational approximation methods. In M. Oppel and D. Saad, editors, *Advanced mean field methods: theory and practice*. MIT Press, 2000.
6. M. Jordan, Z. Ghahramani, T. Jaakkola, and L. Saul. An introduction to variational methods for graphical models. *Machine Learn*, 37:183–233, 1999.
7. D. Knowles and Z. Ghahramani. Infinite sparse factor analysis and infinite independent components analysis. In *ICA*, pages 381–388, 2007.
8. D. Knowles and Z. Ghahramani. Nonparametric bayesian sparse factor models with application to gene expression modelling. *CoRR*, abs/1011.6293, 2010.
9. H. Lee, A. Battle, R. Raina, and A. Ng. Efficient sparse coding algorithms. In *NIPS 22*, pages 801–808, 2007.
10. J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary learning for sparse coding. In *ICML*, page 87, 2009.
11. T. J. Mitchell and J. J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.
12. S. Mohamed, K. Heller, and Z. Ghahramani. Sparse exponential family latent variable models. *NIPS Workshop*, 2010.
13. R. Neal and G. Hinton. A view of the EM algorithm that justifies incremental, sparse, and other variants. In M. I. Jordan, editor, *Learning in Graphical Models*. Kluwer, 1998.
14. B. Olshausen and D. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609, 1996.
15. J. W. Paisley and L. Carin. Nonparametric factor analysis with beta process priors. In *ICML*, page 98, 2009.
16. M. Seeger. Bayesian inference and optimal design for the sparse linear model. *JMLR*, 9:759–813, June 2008.
17. T. Suzuki and M. Sugiyama. Least-squares independent component analysis. *Neural Computation*, 23(1):284–301, 2011.
18. Y. W. Teh, D. Görür, and Z. Ghahramani. Stick-breaking construction for the indian buffet process. *Journal of Machine Learning Research - Proceedings Track*, 2:556–563, 2007.
19. R. Tibshirani. Regression shrinkage and selection via the lasso. *J Royal Statistical Society. Series B*, 58(1):267–288, 1996.
20. M. West. Bayesian factor regression models in the lqlq large p, small n” paradigm. In *Bayesian Statistics 7*, pages 723–732. Oxford University Press, 2003.